

CLAIMS

What is claimed is:

1. A method of content addressable data storage and compression for semi-persistent computer memory for a database management system comprising:
 - 5 providing in the database management system a data structure that associates data identifiers and retrieval keys for memory blocks for storing in semi-persistent memory data from the database management system;
 - storing in the data structure a data identifier;
 - 10 providing a chunk of data comprising a quantity of input data from the database management system;
 - retrieving a memory block from semi-persistent computer memory;
 - 15 searching for a segment of the chunk that matches the memory block; and
 - if a matching segment is found:
 - 20 discarding the matching segment;
 - providing to the database management system a retrieval key for the memory block as a retrieval key for the matching segment;
 - 25 storing in the data structure in the database management system the retrieval key for the matching segment in association with the data identifier;
 - identifying an unmatched portion of the chunk that does not match the

memory block;

30

identifying a free memory block of a file system;

storing the unmatched portion semi-persistently in the free memory block;

35

providing to the database management system a retrieval key for the unmatched portion; and

storing in the data structure in the database management system the retrieval key for the unmatched portion in association with the data identifier.

40

2. The method of claim 1 wherein a free memory block of a file system has a block size at least as large as a maximum memory block size.
3. The method of claim 1 wherein storing the unmatched portion semi-persistently in the free memory block comprises storing the unmatched portion without recording the use of the free memory block in the file system.
4. The method of claim 1 wherein:

identifying a free memory block of a file system comprises reading a block identification from a free block list of a file system, and

5

storing the unmatched portion semi-persistently in the free memory block comprises leaving the block identification unchanged in the free block list of the file system.

5. The method of claim 1 wherein searching for a segment of the chunk that matches the memory block comprises searching at a repeating memory interval through a search section of the chunk for a segment of the chunk that

matches the memory block.

5

6. The method of claim 5 wherein searching at a repeating memory interval through a search section of the chunk for a segment of the chunk that matches the memory block comprises:

5

calculating a weak checksum for the memory block;

calculating weak checksums for segments of the search section of the chunk;

10

comparing the weak checksums for the segments with the checksum for the memory block; and

if a segment is found with a weak checksum equal to the weak checksum of the memory block:

15

calculating a strong checksum for the memory block;

calculating a strong checksum for the segment with the matching weak checksum;

20

comparing the strong checksum of the memory block and the strong checksum for the segment with the equal weak checksum; and

25

determining that the search has found a segment having contents that match the contents of the memory block if the strong checksum of the memory block and the strong checksum for the segment with the matching weak checksum are equal.

7. The method of claim 1 wherein storing the unmatched portion of the chunk comprises storing the unmatched portion of the chunk as a new memory block

having a memory block size equal to the size of the unmatched portion of the chunk.

5

8. The method of claim 1 wherein searching for a segment of the chunk that matches the memory block fails to find a matching segment, the method further comprising repeatedly carrying out the following steps for all memory blocks in computer memory until a matching segment is found:

5

retrieving a next memory block from computer memory; and

searching for a segment of the chunk that matches the next memory block.

9. The method of claim 8 wherein no matching segment is found in any memory block in computer memory, the method further comprising:

identifying a free memory block of a file system;

5

storing a search section of the chunk semi-persistently in the free memory block;

providing to the database management system a retrieval key for the search section of the chunk; and

10

storing the retrieval key for the search section in the data structure in the database management system in association with the data identifier.

10. The method of claim 8 wherein storing a search section of the chunk comprises storing the search section of the chunk as a new memory block having a memory block size equal to the size of the search section of the chunk.

11. The method of claim 8 wherein providing a retrieval key for a search section of a chunk comprises:
- 5 calculating a weak checksum for the search section of the chunk; and
- calculating a strong checksum for the search section of the chunk.
12. The method of claim 1 further comprising:
- receiving a retrieval key from the database management system;
- 5 identifying a memory block in dependence upon the retrieval key;
- retrieving the identified memory block; and
- verifying the contents of the memory block.
- 10 13. The method of claim 12 wherein the retrieval key for the memory block comprises a unique key calculated with an algorithm that generates a unique key from the contents of a memory block, and verifying the contents of the memory block further comprises:
- 5 calculating a new key for the memory block with the same algorithm; and
- comparing the retrieval key and the new key.

14. A system of content addressable data storage and compression for semi-persistent computer memory for a database management system comprising:
- 5 means for providing in the database management system a data structure that associates data identifiers and retrieval keys for memory blocks for storing in semi-persistent memory data from the database management system;
- means for storing in the data structure a data identifier;
- 10 means for providing a chunk of data comprising a quantity of input data from the database management system;
- means for retrieving a memory block from semi-persistent computer memory;
- 15 means for searching for a segment of the chunk that matches the memory block;
- means for discarding a matching segment;
- 20 means for providing to the database management system a retrieval key for the memory block as a retrieval key for the matching segment;
- means for storing in the data structure in the database management system the retrieval key for the matching segment in association with the data identifier;
- 25 means for identifying an unmatched portion of the chunk that does not match the memory block;
- means for identifying a free memory block of a file system;
- 30 means for storing the unmatched portion semi-persistently in the free memory

block;

35 means for providing to the database management system a retrieval key for the unmatched portion; and

means for storing in the data structure in the database management system the retrieval key for the unmatched portion in association with the data identifier.

15. The system of claim 14 wherein:

means for identifying a free memory block of a file system comprises means for reading a block identification from a free block list of a file system, and

5

means for storing the unmatched portion semi-persistently in the free memory block comprises means for leaving the block identification unchanged in the free block list of the file system.

16. The system of claim 14 further comprising:

means for receiving a retrieval key from the database management system;

5 means for identifying a memory block in dependence upon the retrieval key;

means for retrieving the identified memory block; and

means for verifying the contents of the memory block.

10

17. The system of claim 16 wherein the retrieval key for the memory block comprises a unique key calculated with an algorithm that generates a unique key from the contents of a memory block, and means for verifying the contents of the memory block further comprises:

5

means for calculating a new key for the memory block with the same algorithm; and

means for comparing the retrieval key and the new key.

18. A computer program product of content addressable data storage and compression for semi-persistent computer memory for a database management computer program product comprising:
- 5 a recording medium;
- means, recorded on the recording medium, for providing in the database management computer program product a data structure that associates data identifiers and retrieval keys for memory blocks for storing in semi-persistent
- 10 memory data from the database management computer program product;
- means, recorded on the recording medium, for storing in the data structure a data identifier;
- 15 means, recorded on the recording medium, for providing a chunk of data comprising a quantity of input data from the database management computer program product;
- means, recorded on the recording medium, for retrieving a memory block
- 20 from semi-persistent computer memory;
- means, recorded on the recording medium, for searching for a segment of the chunk that matches the memory block;
- 25 means, recorded on the recording medium, for discarding a matching segment;
- means, recorded on the recording medium, for providing to the database management computer program product a retrieval key for the memory block as a retrieval key for the matching segment;
- 30 means, recorded on the recording medium, for storing in the data structure in

the database management computer program product the retrieval key for the matching segment in association with the data identifier;

35 means, recorded on the recording medium, for identifying an unmatched portion of the chunk that does not match the memory block;

means, recorded on the recording medium, for identifying a free memory block of a file computer program product;

40 means, recorded on the recording medium, for storing the unmatched portion semi-persistently in the free memory block;

45 means, recorded on the recording medium, for providing to the database management computer program product a retrieval key for the unmatched portion; and

50 means, recorded on the recording medium, for storing in the data structure in the database management computer program product the retrieval key for the unmatched portion in association with the data identifier.

19. The computer program product of claim 18 further comprising:

means, recorded on the recording medium, for receiving a retrieval key from the database management computer program product;

5 means, recorded on the recording medium, for identifying a memory block in dependence upon the retrieval key;

10 means, recorded on the recording medium, for retrieving the identified memory block; and

means, recorded on the recording medium, for verifying the contents of the memory block.

20. The computer program product of claim 19 wherein the retrieval key for the memory block comprises a unique key calculated with an algorithm that generates a unique key from the contents of a memory block, and means, recorded on the recording medium, for verifying the contents of the memory block further comprises:

means, recorded on the recording medium, for calculating a new key for the memory block with the same algorithm; and

- 10 means, recorded on the recording medium, for comparing the retrieval key and the new key.